



The state of enterprise multi-cloud infrastructure

Cloud service provider performance and the key to realizing the potential of multi-cloud applications.

Summer 2021

Application eXperience Infrastructure Study

Prosimo commissioned Sapio Research, an independent research company, to survey 400 IT experts working in enterprise businesses. We combined this with our own data collected from nearly 500+ PoPs (points of presence) to compare the network transit performance of the three largest cloud service providers (CSPs): Amazon (AWS), Microsoft (Azure) and Google (GCP), worldwide.

Together, these two data sources offer unique insight into enterprises' journey to cloud native and multi-cloud application performance. This information forms the basis of this, the first AXIS report. It provides valuable information for those building and running applications or procuring cloud services.

Contents

Executive summary	01
Applications are vital	02
But performance issues are hampering UX	
Challenges to app experience	04
Apps have left the building	
App diversity is growing	
For user-to-app connectivity, multi-cloud transit wins	08
Security can be a performance killer	14
Diversity impacts performance	19
App diversity is driving the use of multi-cloud	
Data-driven outcomes demand an integrated stack with multi-cloud observability	
Conclusion	26
Methodology	27

Executive summary






Applications are the lifeblood of all enterprises. It is no longer enough for organizations to connect applications and their users. To maintain a competitive advantage, enterprises must deliver fast, reliable application experiences to users in any location on the planet, while implementing Zero Trust frameworks for security.

As such, cloud-native has become the go-to strategy to enable outstanding application experiences at global scale. But as organizations adopt cloud native strategies, they are contending with the complexity of multi-cloud management, the diversity of applications and the global locations of users.

This is the first report to directly assess performance across major cloud providers, as it relates delivering secure application experience. In combination, the report paints a full picture of the enterprise transition to cloud and multi-cloud infrastructure, the challenges and opportunities that cloud architects currently face and the value that multi-cloud infrastructure can provide for global application experience delivery.

While the promise of multi-cloud remains high, enterprises have yet to fully realize its benefits

Our key takeaways from the AXIS research:

<p>1</p>  <p>Enterprises agree that outstanding application experience is key for business success.</p>	<p>2</p>  <p>For user-to-app and app-to-app connectivity, multi-cloud offers a clear path to improving app performance and realizing new efficiencies.</p>	<p>3</p>  <p>Performance varies across cloud providers, regions and applications, signaling the need for dynamic connectivity.</p>	<p>4</p>  <p>Security and complexity remain critical challenges.</p>	<p>5</p>  <p>To realize the promise, enterprises must embrace cloud networking as the first step to help simplify their infrastructure and improve visibility and control.</p>
---	--	--	--	--

Our research makes clear that by gaining maturity in their cloud-native strategies, enterprises will be better positioned to realize new efficiencies, deliver better application experiences and maintain a competitive advantage in contested markets.

Applications are vital

But performance issues are hampering UX


Applications are key — they are the public face of every business. As such, striking a balance between outstanding application experience and efficiency is a business imperative, in order to succeed in today's competitive global marketplace. To better understand this challenge, we polled IT experts on the opportunities and obstacles they face in delivering outstanding application experience today.

Enterprise IT experts are unanimous in their agreement that application experience is vital to business success. For businesses, the benefits of a better application experience are straightforward: improved user and customer satisfaction, increased efficiencies and improved employee productivity. Naturally, application experience is also directly aligned with the top three business priorities for enterprises: increased efficiency, profitability and security. It is unsurprising, then, that expectations around application experience are now exploding.

Today, application experiences must be fast, fluid, secure and ubiquitous anywhere in the world. But it is clear from the research that, with the rising complexity of the infrastructure and applications, organizations will be unable to solve the problem by simply pouring more resources into it.



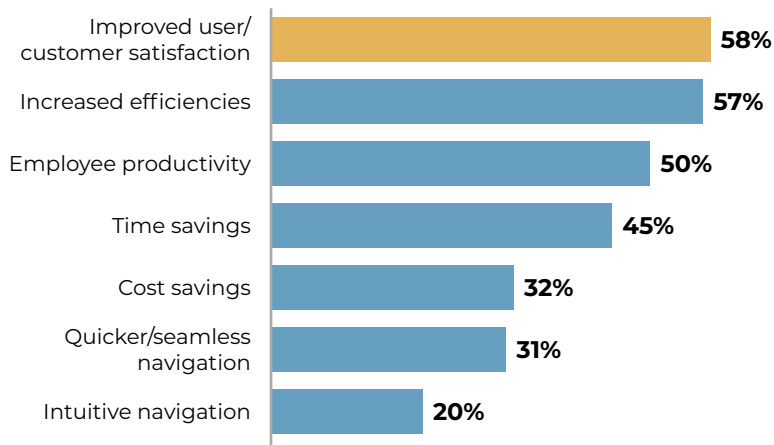
The modern enterprise is focused on profit



The two most common business priorities were increasing efficiency and profit.

Figure 1 What are your organizations top three business priorities for the next 5 years? Select up to three. Base: all respondents [n=400]

Apps are key to satisfaction, efficiency and productivity



Application experience enables productivity and customer satisfaction.

Figure 2 What are the top three benefits of having a great user to application experience? Base: all [n=400]

Key terms used throughout this report

User-to-app connectivity	When a user accesses an application (SaaS or enterprise) located in a private cloud, public cloud or data center.
App-to-app connectivity	When an application, or microservice, communicates with another application or service.
Key routes	Identified connectivity routes between two cloud provider locations (e.g., AWS London to AWS Virginia).
Multi-cloud transit	When traffic is able to traverse the most efficient key routes between two or more cloud service providers.

Enterprises will need both data and operations to enable them to build an effective multi-cloud strategy

Challenges to app experience

Apps have left the building

There are significant challenges in meeting the rising expectations for application experience. These challenges are compounded as users have, literally, left the building. With the Covid-19 pandemic, users have moved from corporate campuses to remote locations, including their homes. While a significant factor, this is just the continuation of the trend for applications to spread out from centralized data centers to public clouds and the edge.

App diversity is growing

At the same time, the diversity of applications is increasing rapidly. Not only are there more kinds of applications, but with the adoption of cloud-native strategies like microservices, applications are being separated from monolithic stacks into dozens of distinct components, each served individually to users.

In combination, these factors have made it a significant challenge to meet the exploding expectations around application experience. Indeed, while 93% of organizations agree that application experience is important for the overall success of the business, roughly two-thirds (64%) find they are unable to offer the same level of experience to everyone, while nearly half (47%) report it is impossible to keep up with the diversity of the application and user landscape.

As a result, enterprises feel not only at greater risk of security incidents, but that they must expend more resources trying to diagnose and resolve application performance challenges like latency, which impacts the majority of large organizations when application experience is poor.

User diversity is changing the business

79% said that users are using a greater variety of devices to access applications.


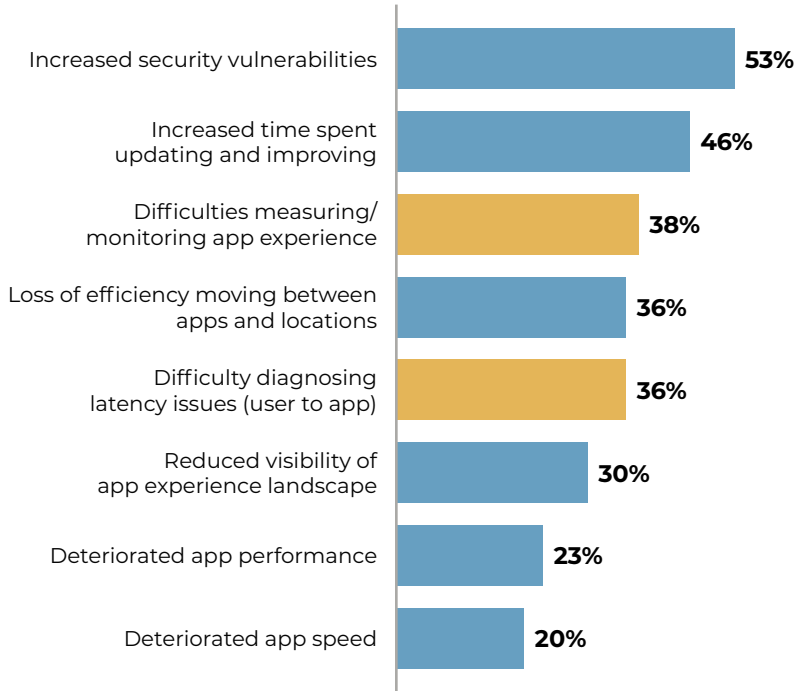


Figure 3 Since lockdown, due to the COVID-19 pandemic and the introduction of widespread home working, how has user diversity within your business changed when it comes to... Base: all respondents [n=400]

This is causing some serious challenges for IT infrastructure

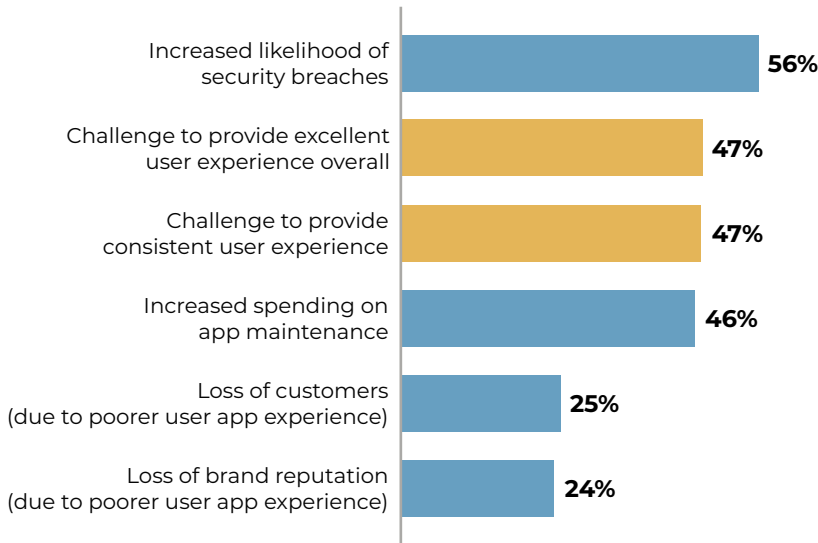


38%

If you can't measure it, you can't improve it. But almost 2 in 5 said they struggled to measure app experience.

Figure 4 What are the consequences of increased user diversity on your application experience? Select all that apply. Base: all respondents [n=400]

Consequences of increased user diversity on the organization



Nearly half (47%) of respondents said they struggle to provide a consistent user experience.

Figure 5 What are the consequences of increased user diversity for your organization? Select all that apply. Base: all respondents [n=400]

93%
 said that app user experience is important to the overall success of the business.

Application diversity directly impacts the applications experience

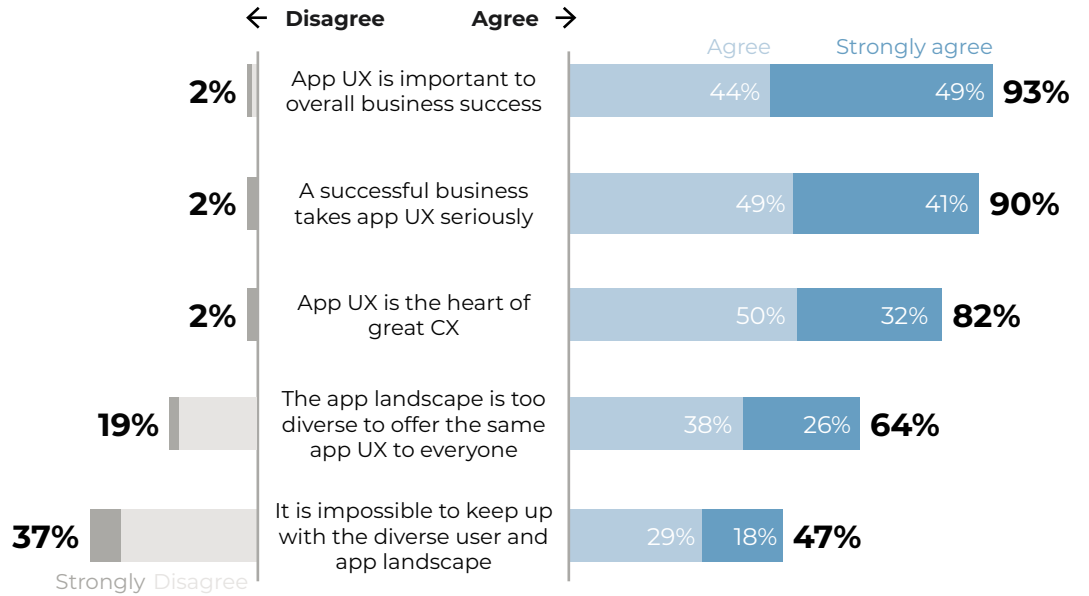


Figure 6 How much do you agree or disagree with the following? Base: all respondents [n=400]

One of the fundamental reasons is network transit – and large organizations with a distributed workforce or those businesses trying to access critical applications are particularly hard hit.

Latency is a common consequence of poor application experience, especially for large companies.

Over half of enterprises report 'latency issues' as a consequence of poor user application experience

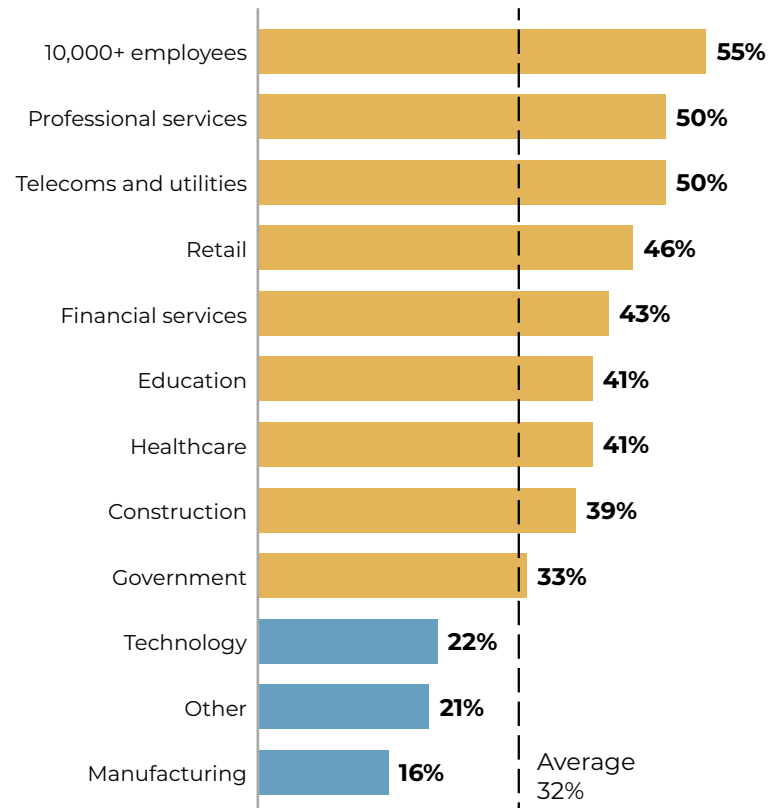



Figure 7 What are the consequences of poor application experience to your organization? Select all that apply.

Base: all respondents [n=400]


Poor application experiences lead to operational challenges



Poor app experience is affecting users, increasing the burden on IT and compromising security.

Figure 8 What are the consequences of poor application experience to your organization? Select all that apply. Base: all respondents [n=400]

Rising complexity results in new challenges



The complexity of applications was the most commonly cited obstacle to delivering excellent app experiences.

Figure 9 What are the main reasons you don't currently offer 'excellent' user application experience? Select all that apply. Note 'Other' = 1%. Base: all those who don't have 'excellent' user experience [n=277]

For many enterprises, the writing is on the wall: they must improve their infrastructure to reliably deliver fast, secure and efficient application experiences. Moreover, with the complexity of managing increasingly diverse applications, enterprises must also simplify their operations and find consistent, centralized ways of delivering application experience at global scale.

For user-to-app connectivity, multi-cloud transit wins

The world's first report comparing global cloud service providers network transit.

For many enterprises, a cloud-native journey that includes the core elements of cloud networking is the logical next step in delivering application experience across a global userbase. The promise of leveraging cloud infrastructure, provided by cloud service providers (CSPs) such as AWS, Azure and GCP, is that enterprises can see improved performance and efficiency at scale.

Located in regions across the globe, CSPs enable enterprises to “place” applications closer to the physical locations of users. CSPs also provide the cloud backbone and cloud networking capabilities to connect users and applications across regions. By leveraging multiple CSPs — a “multi-cloud” strategy — enterprises gain a greater selection of route choices, enabling them to choose the optimum route for a particular user or application. Cloud transit backbones are already used widely to connect applications across regions (app-to-app connectivity), and many enterprises are looking to leverage the cloud backbone for user-to-app access, as well.

We wanted to understand the performance improvements that cloud native strategies can provide organizations today. By collecting latency data from sensors 500+ PoPs worldwide from every major cloud provider, we were able to measure the network transit performance across key routes. We were able to compare the performance of direct forms of connectivity (internet backbone, SD-WAN, VPN, etc.) to single cloud (one CSP) and multi-cloud across regions, where routes are optimized across CSPs. Resulting in the first report to compare network transit across CSPs.

The results are clear. Compared to direct user-to-app connectivity, single cloud provides transit improvements across key routes. Indeed, individual routes can experience more than a 45% improvement in network transit performance. For many enterprises, this fact immediately validates their choice to use single cloud to improve application performance and decrease latency. Still, there is often a variance in latency among CSPs across the same routes. This makes the decision of which CSP to select even harder.

45%
improvement in network transit with multi-cloud.

For user-to-app connectivity, single cloud transit can provide substantial improvements in latency across key routes.

Routes with more than 10% reduction in latency

Figure 10 Routes with more than 10% improvement in latency. Note: this data compares the ingress of users to the closest cloud edge, riding the cloud backbone vs. alternative methods of connectivity (internet, VPN, SD-WAN, backhaul via colo/DCs and the like) to applications hosted in different regions.

Variance across CSPs




Figure 11 Direct vs cloud comparison.

When it comes to overall performance, multi-cloud wins

In terms of network latency performance, neither direct connectivity nor single cloud can match the potential of multi-cloud transit. Not only does multi-cloud offer network latency improvement on 45% more routes than single cloud; it also offers a greater improvement (up to 55%) in network transit, per route when compared against direct methods.

Overall, multi-cloud improves the network transit performance for roughly twice the number of routes, compared to using a single CSP. In other words, the organizations that can master a multi-cloud strategy will wield a true competitive advantage, enabling them to offer faster, more consistent, more efficient application experiences, across more global locations.

Overall, multi-cloud provides substantial user-to-app performance improvements over direct connectivity, across global cloud regions.

Multi-cloud transit impact for user to application traffic

55%
network transit
improvement
per route when
compared against
direct methods.

Figure 12 Key routes that can be improved by multi-cloud for user-to-app.

Enterprises can start with cloud networking, but will need both data and operations to enable them to build an effective multi-cloud strategy.

Grappling with the challenge and opportunity of multi-cloud

To implement a multi-cloud strategy, organizations must do more than just flip a switch. Cloud networking, which encapsulates the core networking requirements for customers looking to interconnect, is a good first step, but it is not enough. That is, in order to optimize pathing across CSPs, enterprises will need both data and operations to support their strategy. For single cloud transit, enterprises must have an understanding of the latency performance data across key routes and CSPs, so they can select the ideal cloud provider for their purposes. This data changes continuously however, so enterprises will need to monitor performance and adapt their strategy on a constant basis. Given that this is the first report to compare network transit data across CSPs and routes, few organizations will be currently able to measure and act on this data.

For multi-cloud transit, this task becomes significantly more complex. In this case, organizations are not only concerned with one cloud, but with (a) monitoring route performance across CSPs and (b) enforcing the optimal path for interconnect between clouds, based on the location of users and best entry and exit point transit for an application. In other words, enterprises need to act on performance data *dynamically* — automating the decision-making process around route choice, so that users and applications experience the greatest possible performance at any given moment.

To manage transit routes across regions and CSPs, most organizations will need to manually log into cloud provider portals, create a VPC, make firewall changes, and more, before making the switch. Then there is the problem of troubleshooting; most enterprises have no reliable way to determine the cause of performance issues, if they are spread across multiple CSPs. Yet, there is variance between routes across multi-cloud — and organizations will need to contend with all of these factors.

There is variance between given routes across multi-cloud. For user-to-app, the most efficient routes are often between regions that are closer together, such as from Korea to Japan in this example from Asia. By optimizing application performance across even shorter, more efficient routes, latencies can be improved significantly. But this requires continuous monitoring and dynamically adapting to optimize the user-to-application access path.

Data is required to determine which routes can be optimized

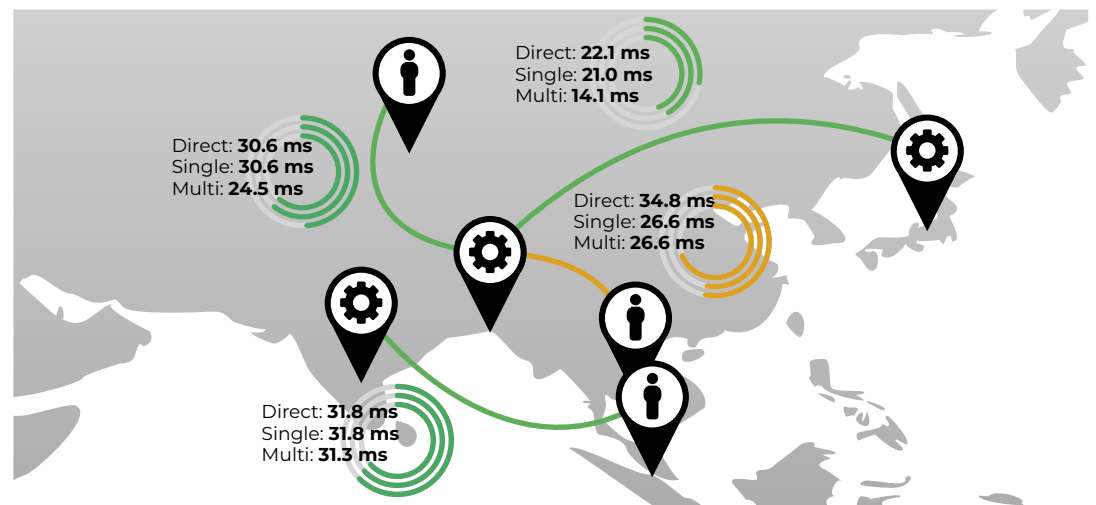


Figure 13 Example of user-to-app routes improved by multi-cloud.

However, this is a leap from what many organizations are capable of today.

☆ Key recommendations

Embrace cloud networking to exploit the strengths of each CSP

To reap the full efficiencies of multi-cloud, enterprises should work to harmonize the power of different cloud providers. Different cloud providers offer different best-in-class services, and their performance varies across different regions of the world. Cloud networking is the first step to interconnect the cloud and it helps by leveraging the strengths of each CSP.

Leverage multi-cloud backbone to improve app experience

Despite the challenges, the promise of multi-cloud remains high, and organizations should still begin using a multi-cloud backbone to deliver better application experience. Even if organizations are not hosting applications in different clouds, such as in both GCP and Azure, the ability to select different routes based on performance can have a positive impact on performance.

Be aware of the necessary trade-offs

Multi-cloud requires cloud architects to balance between cost, security and performance. While CSPs have provided the functionality to scale out network connectivity services by providing native cloud networking capabilities, enterprises should be aware that “over rotating” in any direction comes with trade-offs: a boost to security could come at the expense of performance, reduction in cloud cost can cause performance to suffer. In other words, enterprises must walk a tight-rope and balance competing business objectives.



Security can be a performance killer

Customers need to protect applications and cloud infrastructure.

Security remains one of the highest priorities for application experience — and a persistent challenge. As our survey data indicates, enterprises are more concerned than ever about the growing threat surface caused by cloud transformation and the decentralization of users, applications and data across the globe.

However, with enterprise work-from-home policies of the past year, security has often been put on the backburner, forced as they were to prioritize expanded access and connectivity above all else. Yet, enterprises are also facing a broader range of attacks, particularly to applications accessed through public clouds.

Enterprises are well aware of the security impact caused by rapid cloud transformation over the past 18 months



Figure 14 When it comes to software applications, what are the top 3 benefits of having better security? Select up to three. Base: all respondents [n=400]

It's time to rethink how you protect your users and applications





Figure 15 How much importance do you place on improving the security of applications, compared to how you felt 5 years ago? Note: 0% answered slightly less or much less important. Base: all respondents [n=400]

The range of security threats is also growing. The chart below shows a breakdown of attacks targeting the public cloud applications of just one enterprise over a random 30-day period.

Types of security threats




Enterprises are worried about applications being the new entry point for their IT infrastructure.

Figure 16 Security threats seen over a 30-day period.

There is also the need to secure application access, with limited 1:1 access for users. For the same enterprise, roughly 25% of the traffic in their customer environment is deemed risky, categorized as two users logging in from different locations, multiple login attempts over a short period and the like.

Risk type



Users are still a major risk due to the growth in work from home.

Figure 17 This data shows security logs for a company with 15K users and that has 200 applications, over a 30-day period in April 2021.

Zero Trust Network Access, but at what cost to performance?

Putting security between users and apps hampers performance

Increasingly, enterprise security teams are reacting to this challenge by inserting security vendor SaaS services to enforce user and application security policies across their networks. For many organizations, such SaaS infrastructure provides needed visibility and control over their security posture.

However, this strategy comes with significant tradeoffs. In general, the SaaS model works well for individual app consumption, as evidenced by the skyrocketing levels of SaaS app usage. When it comes to infrastructure, however, SaaS can dramatically increase complexity, due to the need to manage multiple chained SaaS services. Just as importantly, this infrastructure slows performance, as our research shows, caused by the insertion of new services between applications and users. Finally, SaaS infrastructure also creates additional compliance requirements for the enterprise, given that the boundaries of data and privacy straddle infrastructure vendors, SaaS clouds and the organization's own clouds.

Inserting mid-mile SaaS infrastructure can slow performance. In the example below, one application, hosted on a single domain, experienced nine requests to a security services that added a total 342 milliseconds of latency.




Figure 18 Sample data of mid-mile insertion of security between application and user.

Exploring cloud-native security models

Increasingly, enterprises are leveraging models that place security closer to physical location end users — in one example, embedding security in the local edge presence that connects the user to the cloud backbone. This is a similar logic to the cloud-native strategy, in general: reducing latency in the network by bringing services closer to users, while providing additional control over data and the network as the security capabilities reside within the customers control (i.e. in their VPC). In our research, a cloud-native model provides significantly improved performance. Moreover, this model has the benefit of limiting user access to an organization’s cloud applications and network from a single secure, external location.

App response time is 12 ms with native security vs 342 ms.

Bringing security closer to users and the application

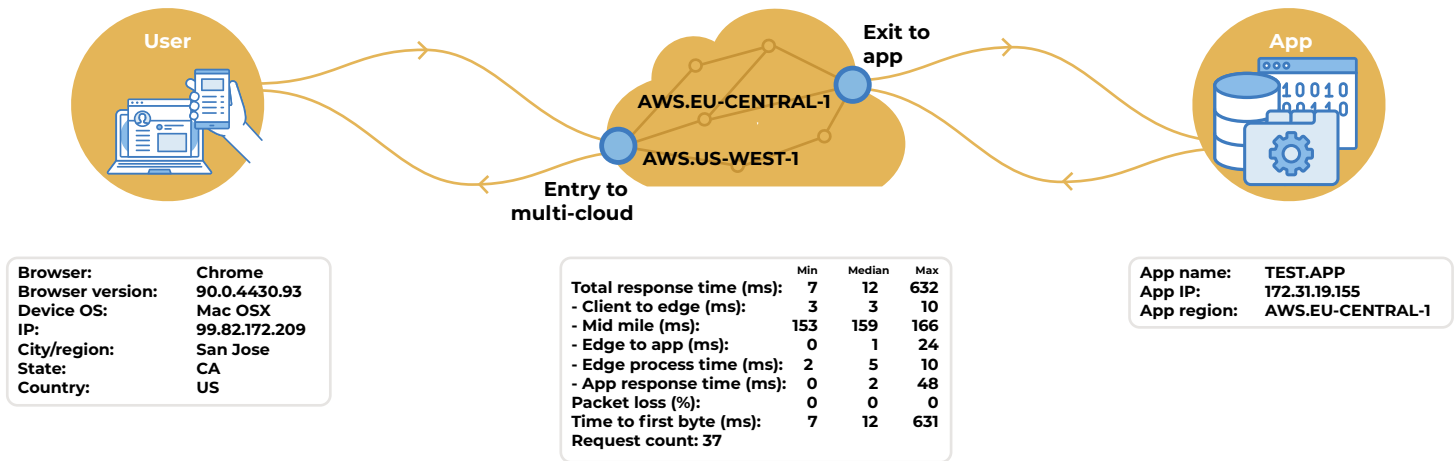


Figure 19 Round-trip measurement of response time with security closer to user and application.

☆ Key recommendations

Security is more important than ever

Moreover, tightly controlling access to cloud applications has become critical, with the global location of users and applications. However, legacy approaches to security like SaaS infrastructure or virtual private networks (VPN) are unlikely to scale at the pace of the cloud, given the increase in complexity and the degradation of performance they entail.

Simpler solutions are critical to implementing ZTNA

If organizations hope to implement Zero Trust frameworks, where all access requires explicit trust (either for user-to-app or app-to-app), enterprises must adopt simpler solutions that match the flexibility of their multi-cloud infrastructure, without sacrificing performance.



Diversity impacts performance

In addition to user-to-app connectivity, we also wanted to understand the multi-cloud opportunity for app-to-app connectivity, given the sheer diversity of applications today. As our research indicates, that diversity is significantly complicating how applications can and should be delivered — creating new challenges for enterprises to optimize application performance across multi-cloud.

Indeed, different types of applications experience significant variations in performance, even when connecting over the same routes, within the same CSP. Our research finds that, overall, the variance can be greater than 15% across applications, even when the same application is accessed from different regions.

Application diversity impacts performance

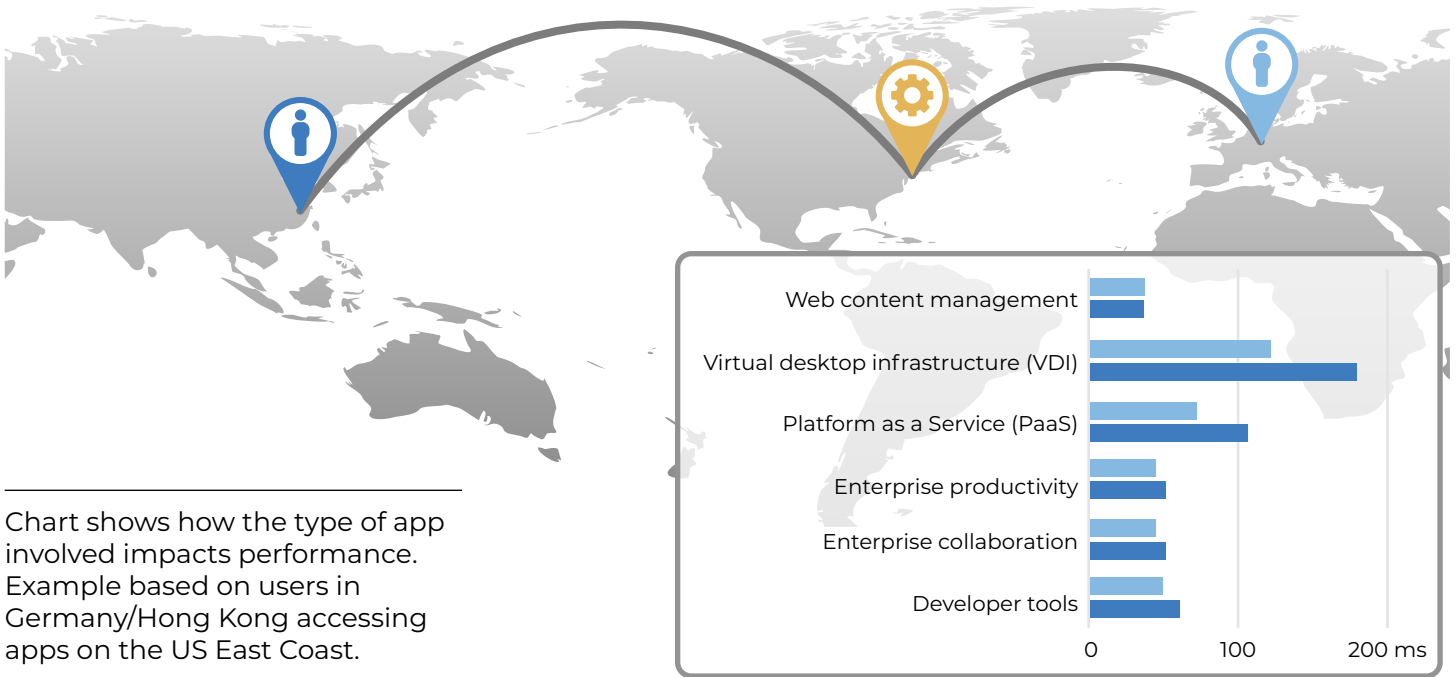


Chart shows how the type of app involved impacts performance. Example based on users in Germany/Hong Kong accessing apps on the US East Coast.

Figure 20 Sample application response time per application category.

A single optimization approach will not work

Indeed, modern applications have vastly different architecture and security requirements. And these differences impact performance, making it difficult for enterprises to monitor and provide universal caching techniques for all applications. This means, in other words, that to optimize performance across multi-cloud, enterprises cannot take a one-size-fits all approach for every type of application. For instance, HTTP applications, native apps (VDI or transaction applications), data lakes, PaaS services, the modern service mesh and more, all have vastly different requirements and need different infrastructure to optimize for experience.

 **Key recommendations****Modern applications have vastly different architectures and requirements**

This diversity of applications makes life difficult for cloud architecture and application delivery teams who want to optimize for performance, as they either do a “lift and shift” of legacy applications or re-create applications in the cloud.

A one-size-fits-all approach won't work

Enterprises need to choose the optimal CSP routes, make the correct caching decisions and dynamically improve performance, based on the type of application.

Companies need the right data

In order to optimize performance for every type of application, enterprises must continuously archive performance data for each individual application type, correlated to performance data of individual CSP routes.

App diversity is driving the use of multi-cloud

With the rise of cloud-native strategies, applications are being split into dozens of microservices, which are then served individually to improve performance and efficiency. This trend is accelerating the use of multi-cloud, given the opportunity to select optimal regions and CSPs, depending on the requirements of microservices. And, when it comes to app-to-app connectivity, selecting the right CSP and multi-cloud route matters. Our research is the first to show that there is significant variance between CSPs across key routes. This impacts app-to-app performance. For example, the chart below shows two routes that experience high variance between providers.

App-to-app latency on key routes

Singapore to the UK (AS-SG to EU-GB)	Azure	AWS	GCP
	100 77.7 ms	+12% 87.1 ms	+62% 125.5 ms
The UK to Hong Kong (EU-GB to AS-HK)	Azure	AWS	GCP
	100 97.1 ms	+6% 102.8 ms	+21% 117.6 ms

Figure 21 App-to-app performance.

The choice of CSP matters. App-to-app performance can vary significantly, depending on route and CSP.

Which CSP is the best for your app-to-app traffic?

Our research is the first to show the significant variance between CSPs across key routes.

Figure 22 Route comparison across CSPs for app-to-app traffic.

Up to
50%
performance
improvement
with multi-cloud
transit for
app-to-app.

For app-to-app, multi-cloud delivers better performance

Given the variance between providers, it is not surprising that multi-cloud transit — selecting the optimum path among all CSP routes — offers consistent performance improvements. Compared to using direct connectivity, multi-cloud transit improves app-to-app performance across roughly 50% of routes, with up to 25% improvement for individual routes. In other words, selecting the ideal CSP in which to place an application component will positively impact the overall performance of the application experience.

Across many routes, multi-cloud delivers more than a 10% improvement compared to direct connectivity.

Impact on app performance with multi-cloud

Figure 23 Routes showing greater than 10% improvement in latency versus direct.

How to choose a route

This chart shows that, for an app located in Brazil connecting with an app located in Singapore, the optimal route is AWS to GCP—if multi-cloud is available. Or, if not, via GCP directly. Not Azure.

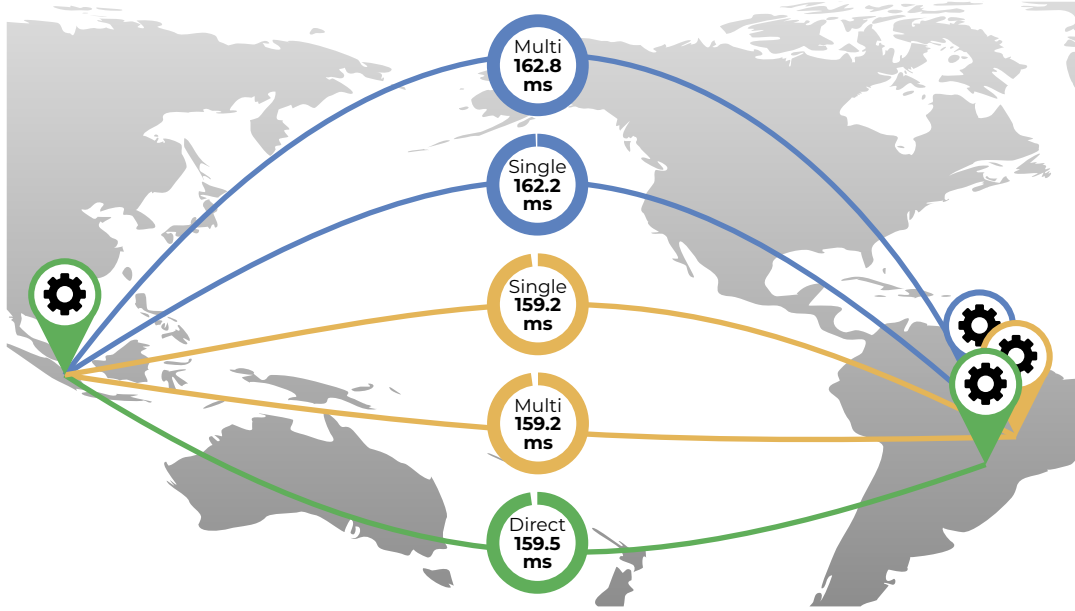


Figure 24 App-to-app example.

☆ Key recommendations

Consider multi-cloud to help improve performance

Given the demonstrable improvements available for application experience, enterprises can no longer rely on traditional optimization approaches — they must consider distributing individual application components across a multi-cloud infrastructure.

Establish closer entry and exit points for users and applications, dynamically

With the variance in performance between cloud providers (and, indeed, between application types) enterprises must contend with the fact that no single optimization stack will realize the full value of multi-cloud. The only sure-fire path to value is to establish closer entry and exit points for users and applications alike. Then, enterprises can gain the ability to provision app-to-app connectivity dynamically, based on the particular applications, routes and CSPs involved.



Data-driven outcomes demand an integrated stack with multi-cloud observability

If enterprises are not currently reaping the full efficiencies of multi-cloud transit, it is not for a lack of measurement. Altogether, 94% of organizations are currently using tools to measure application experience. Yet, measurement does not always equate to true observability. As our research makes clear, enterprises still have significant opportunity to improve the performance of their application experience. And indeed, one the largest obstacles facing enterprise IT professionals is that their tools lack the sophistication to accurately measure application experience.

How you measure app experience impacts performance problem solving

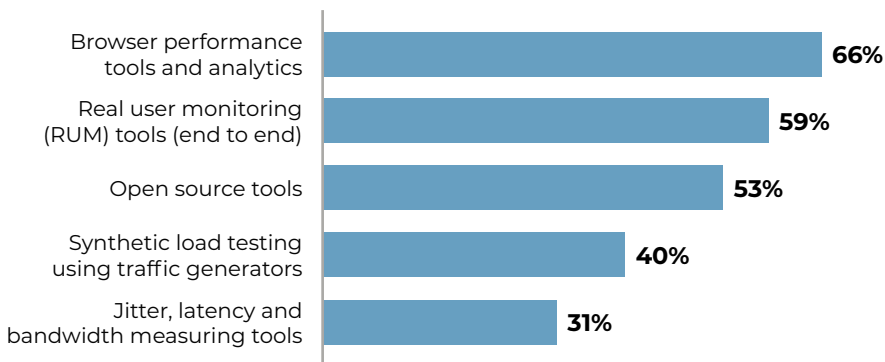


Figure 25 How do you currently measure user application experience performance? Select all that apply. Base: all respondents [n=400]

94%
of businesses currently measure their user application experience performance.

Inadequate measurement tools are a key obstacle to better app experiences

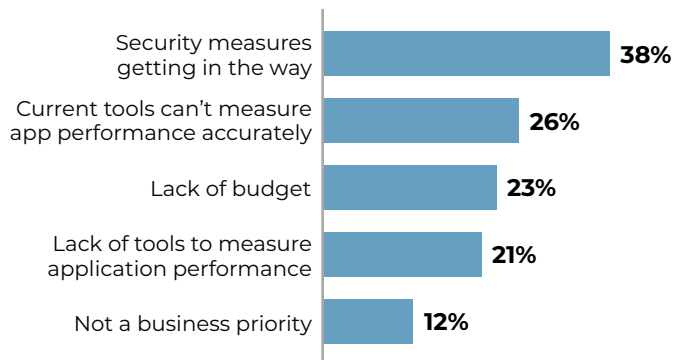


Figure 26 What are the main reasons you don't currently offer 'excellent' user application experience? Select all that apply. Base: all those who don't have 'excellent' user experience [n=277]

The right tools to measure application experience do not exist—existing tools are not up to the job.

Conclusion



Key takeaways

- Infrastructure will change, think simplicity
- Multi-cloud is critical to realizing the benefits of cloud computing on a global scale
- You can improve security by bringing it closer to users and applications
- Leveraging cloud networking can help you provide an excellent app experience consistently globally
- Data will help you improve app experience and avoid wasting money

For many enterprises, the journey to multi-cloud has begun in earnest. On this path, enterprises are poised to deliver faster, more secure application experiences, while improving the efficiency and profitability of the organization. In our twice-per-year AXIS report, we aim to uncover the multi-cloud challenges and opportunities currently faced by organizations on this journey, while arming enterprise IT leaders with the information they need to transform their application experience.

Organizations are frustrated. They need to:

- Manage the evolution of multi-cloud application infrastructure
- Support an increasingly diverse global workforce
- Juggle an ever-changing set of technologies and solutions
- Respond to emerging attack vectors and constantly expanding attack surface
- Mitigate rising cloud costs

For many organizations, 2021 will represent a turning point: it has become clear that traditional infrastructure can't keep up with rising demands for secure access and improved application experience.

As such, it has become critical that organizations unlock the vast potential of the multi-cloud world and support requirements for both user-to-app and app-to-app application experience. While the nuances of multi-cloud delivery are complex, the means to that end are simple: organizations must gain visibility and control across their multi-cloud infrastructure. On the whole, this will allow enterprises to:

- Simplify the architecture
- Get a handle on cloud spend
- Improve security
- Reduce latency and provide consistent experiences anywhere

Looking at the data in this report, many companies will conclude that multi-cloud is the best approach for app delivery. While the improvement to latency is attractive, the real value of multi-cloud ultimately lies in the ability to leverage best-in-class cloud provider capabilities for application performance and security. Of course, once companies pursue a multi-cloud path, operational complexities await. Here, the need for dynamic, modern infrastructure will become even more clear. And this need will intensify as enterprises' multi-cloud strategies develop and reach maturity.

Methodology

1. Real-world performance data

To collect regional inter-cloud latency data, we measure latency between each (supported) cloud region using TCP ping. The data is collected on average every 10 to 15 minutes. Moreover, the cloud regions also ping various servers around the world to measure the public internet latency between cloud regions and geolocations. Using outlier detection techniques, this data is cleaned up to remove any erroneous latency measurements that may be caused by incorrect server locations or local link failures.

2. Survey

Prosimo contracted Sapio Research, an independent research company, to survey 400 US-based IT experts working for businesses with more than 1,000 employees. The interviews were conducted online in Q2 2021. At an overall level, results are accurate to $\pm 4.9\%$ at 95% confidence limits.

Respondents by business size




Figure 27 How many people are employed by your organization? Base: all [n=400]

Respondents by seniority




Figure 29 Which of these best describes your job role? Base: all respondents [n=400]

Respondents by region

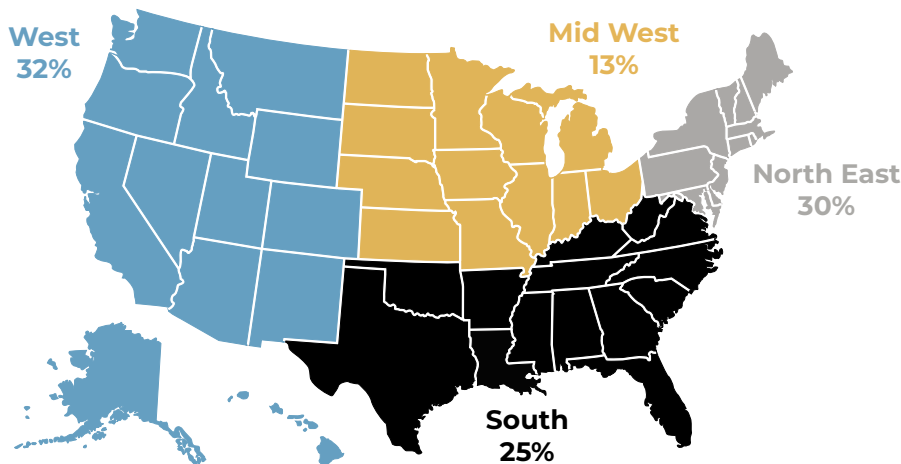


Figure 28 In which region are you based? Select one. Base: all respondents [n=400]

Respondents by function




Figure 30 Which of these departments do you work in? Select all that apply. Base: all [n=400]

Respondents by industry




Figure 31 What sector is your company in? Select one. Base: all respondents [n=400]



Prosimo is a venture-backed company set to disrupt enterprise infrastructure. The company's new, re-imagined architecture provides users in multiple industries with fast and secure experiences for all enterprise applications across a range of environments—quickly, easily, and within a company's administrative control.

To allow customers to benefit from their multi-cloud investments while avoiding operational complexity, Prosimo AXI provides customers with a single vertical stack that brings security closer to users and offers optimized access. Prosimo is built from the ground up with the best that CSPs have to offer. Our infrastructure stack dynamically expands and contracts using rich data insights, self-upgrades, and most importantly, evolves as the CSPs get better services. With Prosimo, customers can successfully scale and transform infrastructure for existing as well as new applications.

Prosimo is backed by General Catalyst and Celesta Capital. The company is headquartered in Santa Clara, California.

For more information, visit prosimo.io.



Sapio Research is a global, full-service market research consultancy. We provide high-quality insights that deliver against key business objectives and inform messaging.

For more information, visit sapioresearch.com.